REMARKS/ARGUMENTS

Request for Continued Examination:

The applicant respectfully requests continued examination of the above-indicated application as per 37 CFR 1.114.

Amendments to all independent claims

Newly added limitations "the JBOD emulation controller is configured to provide a mapping that maps combination of the sections of said group of PSDs to the at least one LMU visible to the host entity, and the at least one LMU is contiguously addressable by the host entity to which the at least one LMU is made available" is merged into all independent claims of the present invention.

Double Patenting

The applicant asserts that claims 1-3, 5-11, 13-44 and 83-109 of the instant application should not be rejected due to double patenting since the cited Nguyen reference (US pat. 7,146,521) fails to teach or suggest the claimed feature of "bringing the LMU on line while the JBOD emulation controller is on line, and taking the LMU off line while the JBOD emulation is on line."

In the Office Action dated December 30, 2009, the Examiner stated that the description in the cited Nguyen reference can be read as the above-mentioned claimed feature, which describes:

> *"removing the overheating disk drive from data storage operations if the controller detects the temperature of the overheating disk drive over a second threshold after the fan speed is increased, wherein the second threshold is greater than the first threshold; and temporarily activating the overheating drive so that data can be evacuated from the overheating disk drive to one or more properly functioning disk drives"*

Nevertheless, the applicant transverses such a rationale, and gives a few points for the difference between the aforementioned feature of the instant application and the teaching of the cited Nguyen reference. These points are as follows:

(a) LMU is not a disk drive

The applicant points out that the claimed LMU is definitely not a disk drive. In fact, a **disk drive** is usually deemed as one type of **physical storage device (PSD)** for those skilled in the art. However, in claim 1, the applicant clearly specifies that "one logical media unit (LMU) comprising **sections of at least one of said group of PSDs**." In other words, the LMU is a virtualized concept, and probably comprises **more than one PSD**. Hence, the applicant believes that the claimed LMU cannot be considered as a single **disk drive**.

(b) "<u>Removing the overheating disk drive" is not "taking the LMU off line"</u>

First of all, according to point (a), the applicant believes that the operation of removing the overheating disk drive is not equal to the operation of taking the LMU off line as the LMU is not a single disk drive. Also, according to Nguyen's teaching (column 7, lines 24-58), it can be found that removing the overheating disk drive is meant to remove **the disk drive that is overheating** from the storage component 18, but not to remove the disk drive that is not overheating, and not to make the RAID off-line. In this way, Nguyen also teaches **reconstructing** the data in the spun down or removed disk drives according to a redundant protocol such as **the RAID-1 or RAID-5** (column 7, line 32); on the other hand, if the RAID has been taken off line, then it is not likely to try to reconstruct the data in the removed one. Thus, it is unreasonable to remove **the whole RAID comprising the overheating disk drive** from the data storage system 10 while the overheating disk drive is removed <u>since such redundant protocols need data in the other disk drives of the **RAID** (e.g. storage component 18) to reconstruct data in the removed one</u>. That is, in the cited Nguyen reference, when one or some of the member disk drive(s) in the same RAID is(are) removed or spun down, the RAID is still on line, and the member disk drives that are properly functioning can still be accessed. However, on the contrary, in the instant application, once the LMU is taken off line, the data in the LMU can not be accessed any more until the LMU is brought on line again.

In addition, in the cited reference, when Nguyen's controller tries to regenerate data in the spun down disk drives but fails, Nguyen teaches **temporarily activating the overheating disk drive,** so that the data can be evacuated from the overheating disk drive into one or more properly functioning disk drive(s) (Nguyen's claim 1). Therefore, the

applicant notes that the cited Nguyen reference **does not take the whole RAID off line but, instead, only removes one or more of the disk drive(s) which is(are) comprised in the RAID**. On the other hand, if Nguyen's teaching were similar to the claimed feature of taking the LMU off line, then no temporary activation operation and data accessing operation would be performed to evacuate the data in the spun down disk drives.

The applicant points out that the operation of taking the LMU off line is meant to take off line **the whole LMU virtualized by sections of the group of PSDs**. Hence, the claimed JBOD emulation controller can never perform data accessing operation on the off-lined LMU once the LMU is taken off line by the JBOD emulation controller.

(c) "Temporarily activating the overheating drive" is not "bringing the LMU on line"

Concerning the operation of temporarily activating the overheating drive, such an operation is meant to temporarily activate the overheating disk drive in order to evacuate the data in the overheating disk drive, and moves data from the overheating disk drive to properly functioning disk drives (see in column 7, lines 42-49 of Nguyen's disclosure). Accordingly, after the data has been evacuated from the overheating disk drive, the overheating disk drive will be de-activated and will not be activated any more. However, the operation of temporarily activating the overheating drive is not equivalent to the operation of bringing LMU on line because (1) the spun down disk drives rather than RAID or LMU are temporarily activated; and (2) no RAID or LMU had been taken off line before temporarily activating the overheating disk drive, and thus, there is no need to bring the RAID or the LMU on line.

On the contrary, the claimed feature of bringing the LMU on line is meant to keep the whole LMU on line, and once the LMU is taken on line, all sections of the group of PSDs comprised in the LMU **can be accessed consistently rather than** temporarily **activated or accessed**.

(d) The objectives of the instant application and the cited Nguyen are different

The objective of the cited Nguyen is "re-constructing the data in the overheating disk drive(s) into other properly functioning disk drives" rather than "taking the LMU off-line when the SVC is on-line and bring the LMU on when the SVC is on- line." Please refer to the abstract and the column 7, lines 30 to line 49 and claim 1 of cited

Nguyen, for such a purpose, there is no need to take the RAID including spun down drives off-line where no data in the RAID can be accessed, and then bring the RAID including spun down drives on-line where the data in the RAID can then be accessed. On the contrary, to achieve such a purpose, it has to keep the RAID always on-line, which includes spun down drives, so that it can try to re-construct the data in the spun down drives from other drives, or to temporarily activate the spun down drives to evacuate the data. If a RAID is off-line, no drive can be accessed and this it is not likely to perform the re-construction of the data, or the temporarily activating and evacuating of the data. Contrarily, **the instant application can selectively take the whole LMU off line or bring the whole LMU on line** under the control of the JBOD emulation controller.

In view of aforementioned points (a) to (d), the applicants asserts that at least the claimed feature of <u>the bringing the LMU on line while the JBOD emulation controller is on line, and taking the LMU off line while the JBOD emulation is on line</u> has never been taught or suggested by the cited Nguyen reference. Thus, claims 1-3, 5-11, 13-44 and 83-109 should be found patentable over claims 1-4, 6, 7, 10-22, 24-27, 29, 30, 32-53, 78-87, 90-93 and 97-103 of copending application No. 10/707,871 in view of the cited Nguyen reference. Withdrawal of the double patenting rejection is respectfully requested.

<u>Claims 1-3, 5-9, 11, 13-17, 24-29, 31-35, 38-40, 44, 83, 84, 86-96, 98-100, 102, 103, 105-107 and 109 are rejected under 35 USC 103a as being unpatentable over Bicknell et al. (US Pub. 2003/0193776) in view of Meehan et al. (US pub. 2004/0177218) and further in view of Nguyen et al. (US pat. 7,146,521).</u>

<u>Independent claims 1, 8, 83, and 92</u>

(A) Regarding newly added limitations "<u>the JBOD emulation controller is configured to provide a mapping that maps combination of the sections of said group of PSDs to the at least one LMU visible to the host entity, and the at least one LMU is contiguously addressable by the host entity to which the at least one LMU is made available</u>" in all amended independent claims 1, 8, 83 and 92

(1a) The logical media unit (LMU) of the present invention is not equal to the

disc pack 118 of disc drive 106 of Bicknell

The LMU of the present invention consists of sections of PSDs, and is formed in a manner that maps combination of the sections of the PSDs and the mapping is from the sections of the PSDs to the LMU visible to the host entity, and the LMU is contiguously addressable by the host entity to which the at least one LMU is made available.

In fact, the LMU is a result after virtualization from the PSDs (such as HDDs). In the present invention, the JBOD emulation controller virtualizes the PSD array into one or more LMUs, to which the JBOD emulation controller makes data access. For example, four HDDs can be virtualized into an LMU of RAID 5, which has a size of three HDDs, or n HDDs can be virtualized into an LMU of RAID 5, which has a size of (n-1) HDDs. Or, four HDDs can be virtualized into a LMU of RAID 1 (mirroring), which has a size of two HDDs, or m HDDs can be virtualized into an LMU of RAID 1, which has a size of (m/2) HDDs. Or, four HDDs can be virtualized into a LMU of RAID 0 (striping), which has a size of four HDDs, or p HDDs can be virtualized into an LMU of RAID 0, which has a size of p HDDs.

In contrast, regarding the meaning of disc pack 118 in figure 3 of Bicknell, the disc pack 118 is a stack of discs mounted on a spindle motor for co-rotation about central axis 122. Therefore, **the disc pack 118 is a component part in the disc drive 106 (HDD 106), and the disc drive 106 is a kind of PSD.**

Please refer to the following paragraph 0018 and fig. 3 of Bicknell.

FIG. 3

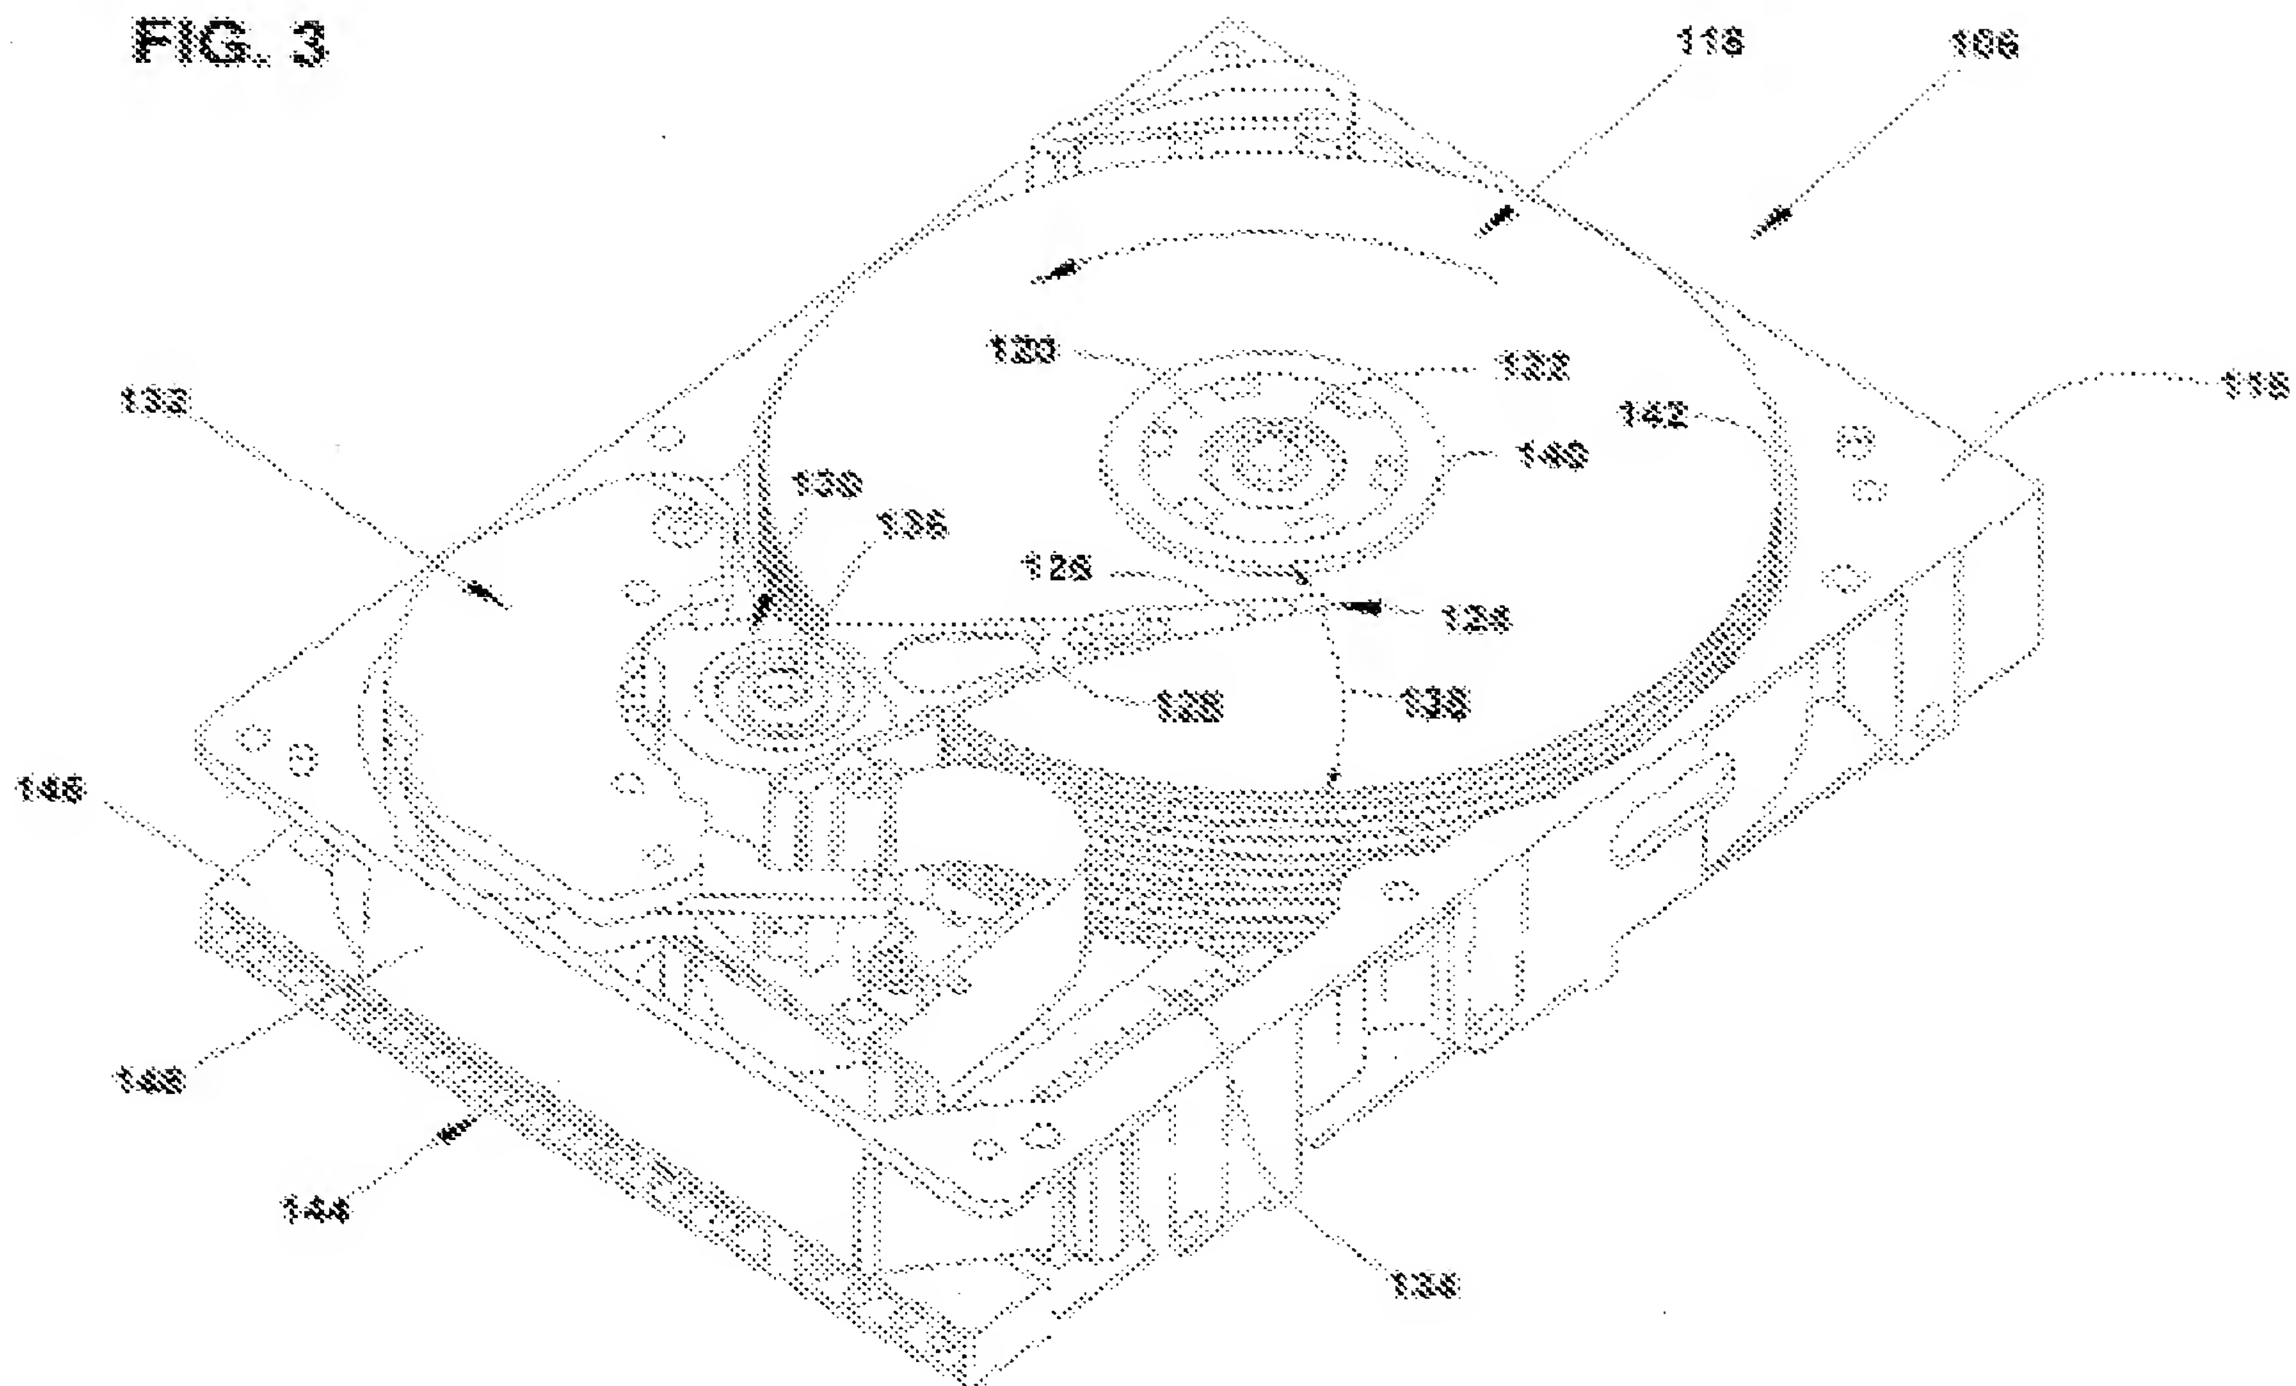

[0018] Disc drive assemblies 104 each include a disc drive 106 that is contained in a carrier 114. FIG. 3 is an isometric view of an example of a disc drive 106 that can be used with disc storage subsystem 100. Disc drive 106 includes a housing with a base 116 and a top cover (not shown). Disc *drive 106 also includes a disc pack 118, which is mounted on a spindle motor (not shown) by a disc clamp 120. Disc pack 118 includes a plurality of individual discs, which are mounted for co-rotation about central axis 122. Each disc surface has an associated disc head slider 124 which is mounted to disc drive 106 for communication with the disc surface. In the example shown in FIG. 3, sliders 124 are supported by suspensions 126,*

> *which are in turn attached to track accessing arms 128 of an actuator 130. Actuator 130, shown in FIG. 3, is of the type known as a rotary moving coil actuator and includes a voice coil motor, shown generally at 132. Voice coil motor 132, under the control of servo electronics 134, rotates actuator 130 with its attached heads 124 about a pivot shaft 136 to position heads 124 over a desired track along an arcuate path 138 between a disc inner diameter 140 and a disc outer diameter 142.*

The above figure 3 and paragraph 0018 of Bicknell disclose the structure of a hard disc drive 106 (i.e., disc drive 106). "The disc drive 106 includes **a housing with a base 116 and a top cover (not shown).**" "Disc drive 106 also includes a disc pack 118, which is mounted on a spindle motor (not shown) by a disc clamp 120. Disc pack 118 includes a plurality of individual discs, which are mounted for co-rotation about central axis 122."

That is, the disc in the disc pack 118 is the storage media having magnetic material coated thereon for storing data therein. Therefore, **the disc is a component in the disc drive 106.**

Moreover, the disc pack 118 is a stack of discs mounted on a spindle motor for co-rotation about central axis 122. Therefore, **the disc pack 118 is also a component in the disc drive 106.**

Furthermore, the "**housing with a base 116 and a top cover" is simple a casing for a disc drive 106**, not an enclosure for a plurality of HDDs.

From the above explanations to the LMU of the present invention and the figure 3 and paragraph 0018 of Bicknell, it can be concluded that **disc pack 118 is a component part in a hard disc drive 106**, and thus is totally different from the LMU of the present invention, which is formed by the fact **that the SVC virtualizes a plurality of PSDs into one or more LMUs visible to and contiguously addressable**

<u>by the host entity.</u>

(1b) The logical media unit (LMU) of the present invention is not equal to the data interface 144 of Fig.6 of Bicknell

Regarding meaning of LMU, please refer to aforesaid point (1a) that has already explained the meaning of LMU of the present invention.

In contrast, please refer to the following Fig.6 and paragraph [0019], in which "Disc drive 106 also includes <u>a data interface 144</u> including a connector 146 located at rear side 148, through which data is communicated and power is provided to disc drive 106. <u>Data interface 144 is a standardized data interface, such as either a Serial Advanced Technology Architecture (SATA) interface or a Parallel Advanced Technology Architecture (PATA) interface.</u> Other data interfaces are also possible for disc drive 106."

From above explanations, it can be concluded that the data interface 144 of Fig.6 of Bicknell is an standardized data interface like SATA interface or PATA interface, and thus is totally different from the LMU of the present invention, which is formed by the fact **<u>that the SVC virtualizes a plurality of PSDs into one or more LMUs visible to and contiguously addressable by the host entity.</u>**

FIG. 6

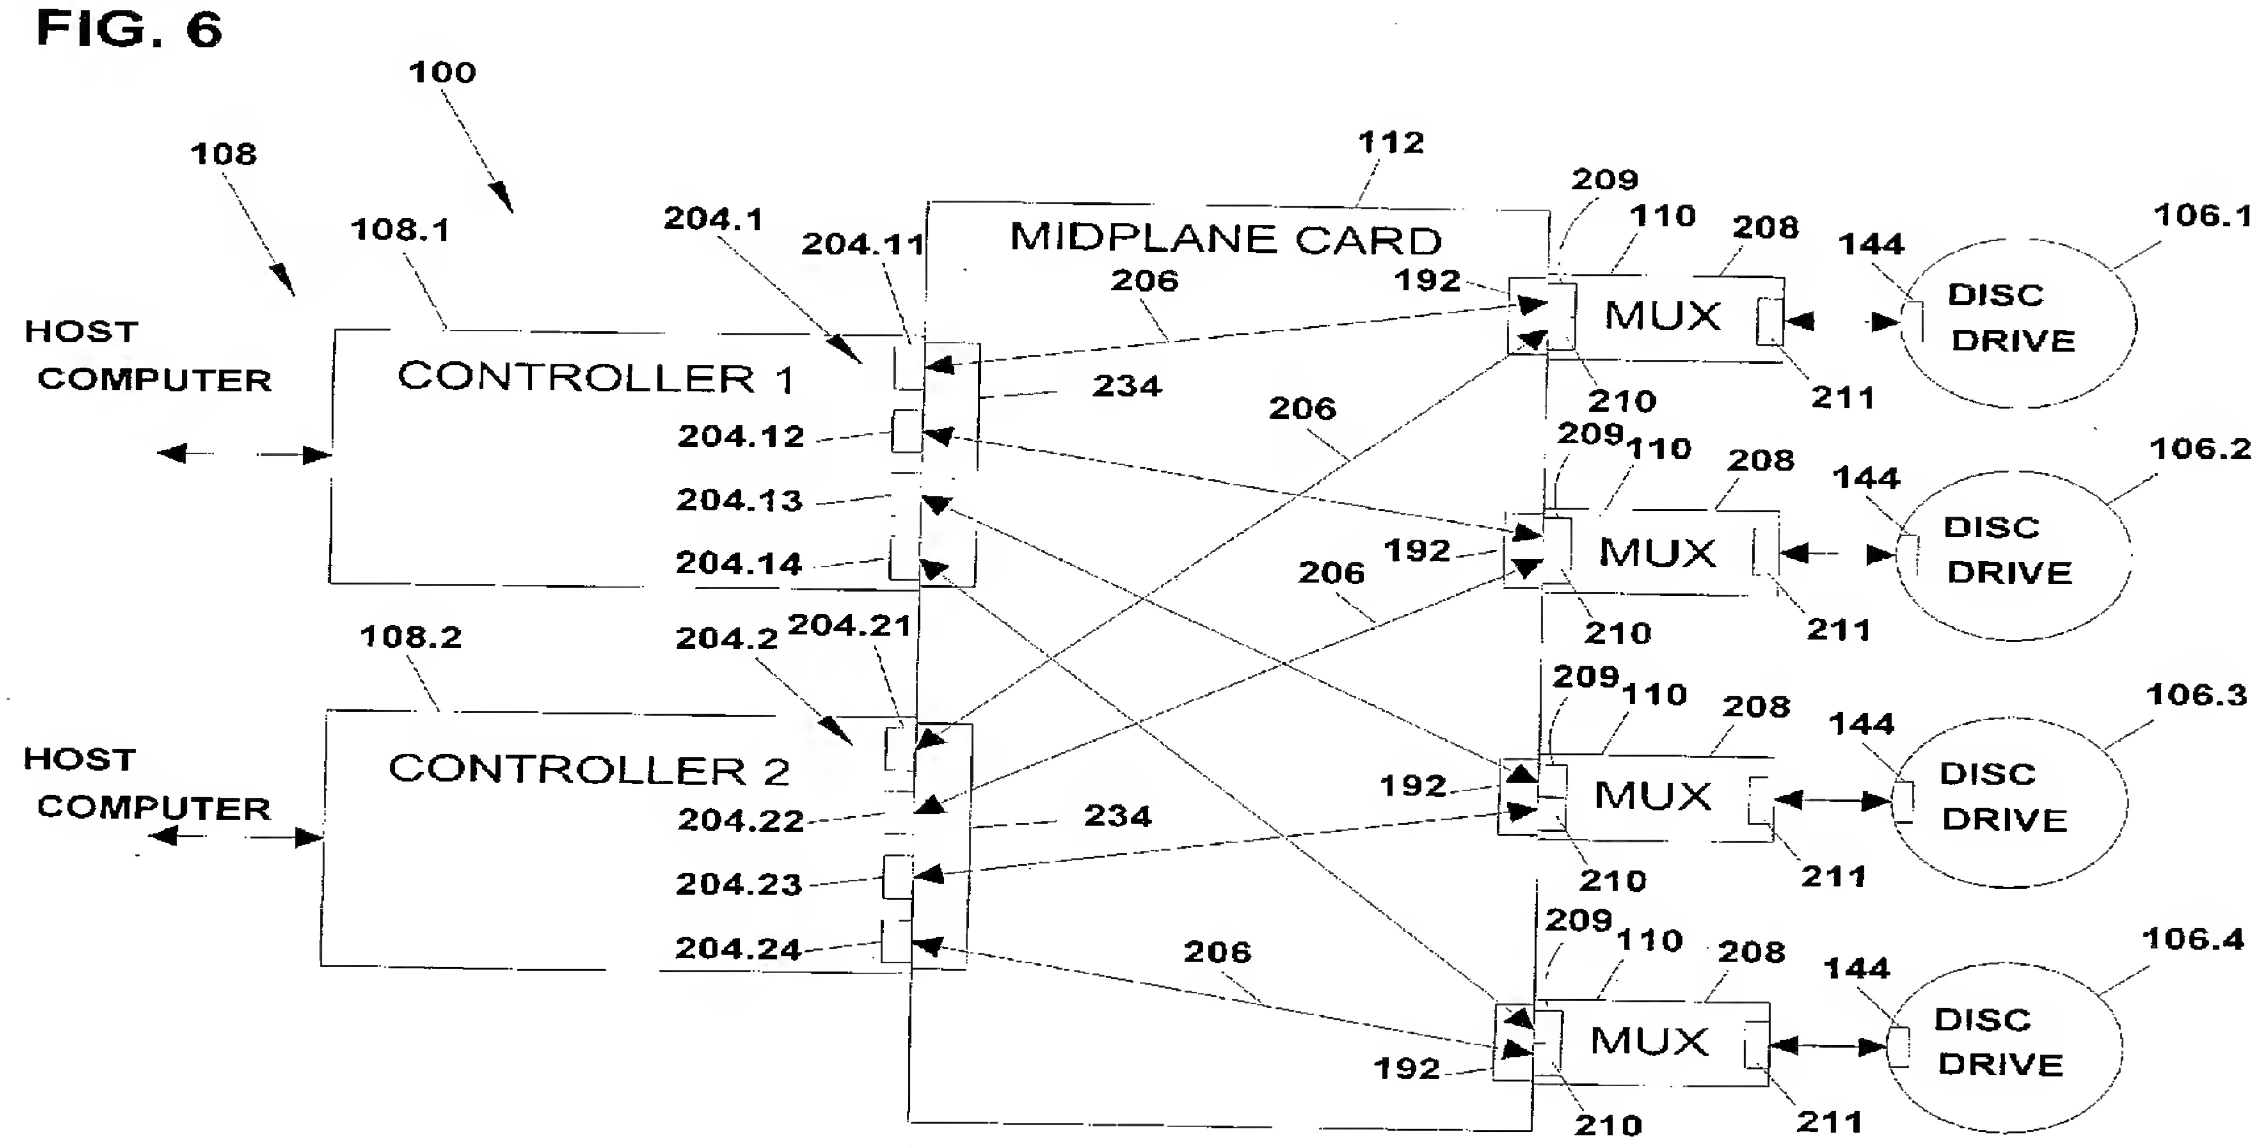

[0019] Disc drive 106 also includes a data interface 144 including a connector 146 located at rear side 148, through which data is communicated and power is provided to disc drive 106. Data interface 144 is a standardized data interface, such as either a Serial Advanced Technology Architecture (SATA) interface or a Parallel Advanced Technology Architecture (PATA) interface. Other data interfaces are also possible for disc drive 106. Disc drive 106 can preferably be removed without disturbing the operation of subsystem 100. Additionally, a Redundant Array of Independent Disks (RAID) architecture is preferably employed in disc storage subsystem 100 where subsystem 100 includes one or more additional disc drives 106 storing redundant data. This arrangement provides security against data loss due to disc drive failure

(B) Regarding <u>"bringing the LMU on line while the JBOD emulation controller is on line, and taking the LMU off line while the JBOD emulation controller is on line"</u> in all amended independent claims 1, 8, 83, and 92

Concerning the patentability of claim 1, the applicant asserts that the claimed feature, <u>"bringing the LMU on line while the JBOD emulation controller is on line, and taking the LMU off line while the JBOD emulation controller is on line"</u> has never been taught or suggested by the cited Nguyen reference based upon the same argument made to the double patenting rejection of claim 1. Thus, the applicant asserts that claim 1 of the present invention should be found allowable in view of the cited references because the combination of the cited reference fails to teach all of the features recited in claim 1. Similar arguments also apply to independent claims 8, 83, and 92; and the various dependent claims should also be found allowable for at least the same reasons as their respective base claims. Consideration of claims 1-3, 5-9, 11, 13-17, 24-29, 31-35, 38-40, 44, 83, 84, 86-96, 98-100, 102, 103, 105-107 and 109 in view of these arguments is respectfully requested.

Further comments regarding the patentability of particular dependent claims with respect to the cited references are provided as below.

<u>Dependent claims 13, 86 and 89</u>

The instant application claims comprising <u>**auto-on-lining mechanism**</u> to automatically bring on line a said LMU which was previously off-line once <u>**a requisite quorum of said PSDs**</u> comes on-line, the meaning of which is explained in paragraphs 0076 and 0077. Please refer to [0076] and [0077] of the present invention.

> *Paragraph [0076] "An important enabling element of the automated on-lining and off-lining of LMUs in response to PSD insertions/removals is a facility for detecting if a PSD 36 is inserted/removed into/from the JBOD enclosure. FIG. 15 shows an*

example of such a detection facility. The JBOD emulation controller 38 monitors the state of this detection facility to determine when a PSD 36 is removed or inserted. In the simplest case, in which each logical media unit is composed of a single PSD, removal of the PSD will trigger the initiation of the procedure that takes the corresponding logical media unit off line, while insertion will trigger the procedure to scan in and then bring on line the corresponding logical media unit."

[0077] In order to support the above-mentioned features of the present invention JBOD subsystem, the JBOD subsystem can further comprises an auto-on-lining mechanism to automatically bring on line a said logical media unit which was previously off-line once a requisite quorum of said PSDs comes on-line, an auto-off-lining mechanism to automatically take off line a said logical media unit which was previously on-line once a requisite quorum of said PSDs becomes off-line, a determining mechanism for automatically determining when a PSD has been removed or when one has been inserted, and a scanning-in mechanism to automatically scan in PSDs on detection of insertion of the PSD.

In contrast, as described above, Nguyen never teaches anything regarding taking a whole RAID (e.g. LMU) off-line, so it is impossible that Nguyen discloses an auto-on-lining mechanism to automatically bring on line a said LMU which was previously off-line. Also, note is respectfully made by the applicant that Nguyen merely teaches **temporarily activating an overheating disk drive rather than bringing a whole RAID on line.**

<u>Dependent claims 14, 87 and 90</u>

The instant application claims <u>"further comprising auto-off-lining mechanism to automatically take off line one of said LMU which was previously on-line once a</u>

requisite quorum of said PSDs becomes off-line."

Similar to the argument made to claim 13, please refer to arguments regarding aforesaid claim 13 of the present invention. Nguyen only teaches removing an overheating disk drive rather than taking off line a whole RAID. Thus, the claimed feature of **auto-off-lining mechanism** to automatically take off line a said LMU is by no means taught or suggested by Nguyen.

Dependent claim 16

The instant application claims "comprising scanning-in mechanism to automatically scan in PSDs on detection of insertion of the PSD." Please refer to [0077] of the present invention in which [0077] in order to support the above-mentioned features of the present invention JBOD subsystem, the JBOD subsystem can further comprises an auto-on-lining mechanism to automatically bring on line a said logical media unit which was previously off-line once a requisite quorum of said PSDs comes on-line, an auto-off-lining mechanism to automatically take off line a said logical media unit which was previously on-line once a requisite quorum of said PSDs becomes off-line, a determining mechanism for automatically determining when a PSD has been removed or when one has been inserted, and a scanning-in mechanism to automatically scan in PSDs on detection of insertion of the PSD. That is, the scanning-in mechanism of the present invention automatically scan in PSDs upon detection of insertion of the PSD.

In contrast, please refer to paragraph 0014 of Meehan (see below), in which "the host interface 202 may comprise a "plug-in" card that is inserted into a backplane of a computer system (e.g., server), and the Primary RAID Controller 205 may communicate with this host interface card via a cable.", but obviously fails to teach or suggest "such a scanning-in mechanism" of the present invention at all, not to speak of "automatically scanning in PSDs on detection of insertion of the PSD" of the present invention.

[0014] FIG. 3 illustrates a block diagram of a RAID

architecture 200, according to one embodiment of the present disclosure. Referring to FIG. 3, the RAID architecture 200 includes a primary RAID controller 205 at a first RAID level (or stage) and "m" secondary RAID controllers 210 (nodes) at a secondary RAID level (or stage), where "m" is a positive whole number greater than one. The RAID architecture 200 is typically implemented in conjunction with a computer system (not shown) where the RAID controller 205 communicates with (by writing data to and reading data from the storage disks 230) a central processing unit or other component(s) of the computer system via the host interface 202. For example, <u>the host interface 202 may comprise a "plug-in" card that is inserted into a backplane of a computer system (e.g., server), and the Primary RAID Controller 205 may communicate with this host interface card via a cable.</u> By way of another example, the Primary RAID Controller 205 may be implemented on the "plug-in" card or on a motherboard of the computer system, and is coupled to the Secondary RAID Controllers 210 via a communication medium (e.g., cable).

<u>Dependent claim 17</u>

The instant application claims <u>"comprising informing mechanism for informing the host entity when the mapping of said LMUs to host-side interconnect LUNs has changed."</u> Please refer to paragraph 0079 of the present invention, in which <u>[0079] In general, when the mapping of LMUs to the host-side IO device interconnect LUNs has changed, either when an LMU is newly introduced onto the Fibre loop or when an existing LMU is removed therefrom, the host entity will be informed of such change.</u> An example of such JBOD emulation controller is that when the host-side IO device interconnect is a Fibre operating in Arbitrated Loop mode, while the external JBOD emulation controller can issue a LIP when a new target ID is introduced onto the Fibre loop or is removed from the Fibre loop so as to inform other devices on the loop that the loop device map has changed. Another example is that when the JBOD emulation

controller support standard SCSI command set as a primary command interface with the host entity over the host-side IO device interconnect, the JBOD emulation controller posts a status of "CHECK CONDITION" to the host entity with sense key of "UNIT ATTENTION" and sense code of "REPORTED LUNS DATA HAS CHANGED" to inform the host entity when the mapping of LMUs to the host-side IO device interconnect LUNs has changed. That is, for example, when an LMU is newly introduced onto the Fibre loop or when an existing LMU is removed therefrom, i.e., the mapping of LMUs to the host-side IO device interconnect LUNs has changed, the host entity will be informed of such a change of mapping of LMUs to the host-side interconnect LUNs. In other words, change of mapping depends on insertion/removal of the LMU. Upon insertion/removal of the LMU, change of mapping to the host-side IO device interconnect LUNs happens.

In contrast, paragraph 0017 of Bicknell only discloses "Each intermediate electronic component 110 determines which controller 108 is provided data access to a particular disc drive 106 by opening and closing data communication paths between the disc drive 106 and each of the controller 108. In the event that one of the controller 108 fails, data stored in the disc drives 106 can still be accessed by the host computer through the remaining active controller 108. In this manner, the reliability of disc storage subsystem 100 is improved", but fails to teach or suggest <u>"informing the host entity when the mapping of said LMUs to host-side interconnect LUNs has changed" of the present invention."</u>

<u>Dependent claim 32</u>

Claim 32 of the instant application claim "comprising an enclosure management services mechanism." Please refer to [0093] of the instant application in which "Yet another feature of a JBOD subsystem is enclosure management services (EMS). This is an intelligent circuitry **<u>that monitors status of various enclosure devices, such as power supplies, fans, temperatures, etc. and can be interrogated by a host for these statuses.</u>**" That is, the instant application claims comprising an enclosure management services (EMS) mechanism, in which the EMS mechanism is

an intelligent circuitry that monitors status of various enclosure devices, such as power supplies, fans, temperatures, etc. and can be interrogated by a host for these statuses.

In contrast, paragraph 0037 of Bicknell only discloses "the multiplexing electronics selectively opens and closes the first and second data communication paths in response to at least one control signal (such as 218 or 220)". That is, **MUX 208 of Bicknell** discloses a "**switching electronics"** to selectively open one of the first and second data communication paths and close the other of the two data communication paths, which is **"not power supplies."** Mux can not perform the functions the EMS mechanism like **monitoring status of various enclosure devices, such as power supplies, fans, temperatures, etc. power supplies, fans, temperatures, etc.** In other words, **obviously, MUX 208 of Bicknell is not equal to the EMS of the present invention, and naturally** fails to disclose, such as the **EMS mechanism which is an intelligent circuitry that monitors status of various enclosure devices, such as power supplies, fans, temperatures, etc. power supplies, fans, temperatures, etc.**

Conclusion

Thus, all pending claims are submitted to be in condition for allowance with respect to the cited art for at least the reasons presented above. The Examiner is encouraged to telephone the undersigned if there are informalities that can be resolved in a phone conversation, or if the Examiner has any ideas or suggestions for further advancing the prosecution of this case.

Recognizing that Internet communications are not secure, I hereby authorize the USPTO to communicate with me concerning any subject matter of this application by electronic mail. I understand that a copy of these communications will be made of record in the application file.

Sincerely yours,

/Winston Hsu/ Date: 06/30/2010

Winston Hsu, Patent Agent No. 41,526
P.O. BOX 506, Merrifield, VA 22116, U.S.A.
Voice Mail: 302-729-1562
Facsimile: 806-498-6673
e-mail : winstonhsu@naipo.com

Note: Please leave a message in my voice mail if you need to talk to me. (The time in D.C. is 12 hours behind the Taiwan time, i.e. 9 AM in D.C. = 9 PM in Taiwan.)